

**The Long Hold: Storing Data at the National Archives**

Kenneth Thibodeau, Ph. D.  
Director, Center for Electronic Records  
National Archives and Records Administration

The National Archives is, in many respects, in a unique position. For example, I find people from other organizations describing an archival medium as one which will last for three to five years. At the National Archives, we deal with the centuries, not years. From our perspective, there is no archival medium for data storage, and we do not expect there will ever be one. Predicting the long-term future of information technology, beyond a mere five or ten years, approaches the occult arts. But one prediction is probably safe. It is that the technology will continue to change, at least until analysts start talking about the post-information age. If we did have a medium which lasted a hundred years or longer, we probably would not have a device capable of reading it.

The issue of obsolescence, as opposed to media stability, is more complex and more costly. It is especially complex at the National Archives because of two other aspects of our peculiar position. The first aspect is that we deal with incoherent data. The second is that we are charged with satisfying unknown and unknowable requirements.

The data is incoherent because it comes from a wide range of independent sources; it covers unrelated subjects; and it is organized and encoded in ways that not only do we not control but often we do not know until we receive the data.

The sources are potentially any operation of the Federal Government, or its contractors. The National Archives has been in the business of collecting digital data for two decades. The way we get it is through our authority over all federal records. Under the Federal Records Act, no agency of the Federal Government can destroy or alienate any Federal record without authorization from the Archivist of the United States, who is the head of the National Archives and Records Administration. Simplistically, the way it works is that agencies tell us what records they have, and we tell them which ones they can destroy when they no longer need them, and which ones must be preserved for posterity. (The definition of Federal record in the law explicitly includes machine-readable files.)

Since 1972, we have reached agreements with agencies that provide for them to transfer to us, and for us to preserve, data from 600 data collections. 573 of these are still active. From these agreements, we have received over 10,000 data files. The rate of transfer has increased dramatically in the last two years: In fiscal year 1988, the National Archives received 167 data files. So we are currently operating at eight times the volume of new files we had years ago, and we expect at least to double next year.

Those numbers are very encouraging, but the overall picture is rather bleak. If we look at all of the data which was scheduled to arrive in the last twenty years, from those 600 data collections, we have received less than 7% of the transfers which should have been made. We recently completed development of a system to generate dunning letters to agencies who fail to transfer data as scheduled, and to track each case to completion. But this system creates additional problems. If I implement it as planned, on a government wide basis, we would need to increase our capability to handle new files, not by doubling current capacity, but by increasing it more than six times. And to handle the backlog of data which should have come in before now, I would need at least 10 times our current capacity.

The past gives us pause. But the future is a brave new world. At least it requires a degree of bravura just to glance in that direction. We have underway a study which is looking beyond the 600 data collections we have decided to preserve to see what else is out there. It is a study of major federal databases being conducted by the National Academy of Public Administration

(NAPA). This study has some interesting exclusions. First of all, we told NAPA not to bother with systems used for generic housekeeping functions, such as personnel, payroll, procurement and supply, because there is little likelihood that we would have any interest in preserving data from such a system. Secondly, we told them not to look at big science, because that is such a large and complex area that it deserves separate attention. (We hope to engage in a project with the National Academy of Sciences on the preservation of scientific data.) Thirdly, we told NAPA not to worry too much about databases on PCs, simply because they would never finish the project if they tried to find all the interesting databases sitting on desktops. With those limitations, NAPA has identified over 10,000 databases.

Obviously, that is far too big a number even for us to think about. So we gave NAPA a set of criteria for culling from the total inventory a subset of those databases with some likelihood that the National Archives would be interested in preserving them. We thought we might wind up with a list of the 500 most important databases in the Federal Government, from an archival perspective. That list would pose quite a challenge for us, because it could practically double the total number of data collections generating data that we want to preserve. The subset of 500 currently has about 900 members.

The next phase of this study is to solicit advice from subject area experts about what data we should try to preserve. NAPA has organized five working groups, with a total of 32 experts in a variety of fields. We are bringing these people together at the end of July for a four day meeting where they will try to develop some common opinions on the long term value of the data.

Which brings me back to the basic point here: what we are dealing with is incoherent data. It concerns practically any area in which the United States Government is involved, which is practically anything. The data we already have ranges from data about tektites on the ocean floor to military operations in time of war. It includes census data on population and the economy, data on Japanese-American internees in World War II, detailed data on air traffic and on stock and bond transactions, and on many, many other subjects. The variety of subjects covered is also increasing.

The data is extremely diverse in content, but content is often the only thing we know about the data until it comes in. We know how many transfers are due, but most often we do not know what the volume of data in a transfer will be, or how it will be organized, even at the physical file level. For example, the files which came in during the first six months of this fiscal year ranged in size from 6 K to 1.4 gigabytes. The number of files in a transfer has ranged from one to 400, and we expect some transfers in the next few years will contain thousands of files.

One thing we do know about the data before it arrives is its logical structure: everything we receive is in flat file format, because we require it to come in that form. However, we realize that this requirement is unreasonable and unrealistic in many cases. We are working to expand the range of formats we will accept to include relational tables. We expect to change our regulation to that effect by the end of this year. We know that, when we do that, it will be only one of many steps we will have to take in a journey with no foreseeable end.

That is a brief overview of one aspect of the unique situation of the National Archives. The second aspect is that we are charged with satisfying unknown and unknowable requirements.

NARA's mission to preserve and provide access to records with enduring value makes NARA, in effect, the agent of generations yet unborn. What differentiates this agency from other parts of the government is the unique responsibility NARA has to serve the information needs of the distant future. This responsibility is fundamental to the very essence of the National Archives as keeper of the Nation's memory.

NARA's responsibility to the future places us in a perpetual quandary: we must devote ourselves to serving needs which we cannot know. We cannot know the questions the future will ask of its past, nor how future researchers will go about answering these questions. We must assume, however, that the information technology which will be available in the future --

- even in the very near future --- will be more powerful and more flexible than what is available today. Information processing problems which today are difficult and costly, if not impossible to solve will become as simple as getting a computer to print out narrative in paragraph form. (A short 20 years ago that was beyond state of the art.)

Along with the technology, analytic tools will continue to improve: there will be further developments as powerful as the mathematics of chaos which will help researchers to understand things which today appear to defy reason. We can also assume that events will happen in the future, which will be as threatening as the depletion of atmospheric ozone, or as exciting as Operation Desert Storm, or as commonplace as the passing of generations, which will make future users want to go back to reexamine the records of the past.

EVENING SESSION

(8:30 p.m.)

DR. HARIHARAN: Ladies and gentlemen, it is my pleasure again to introduce Dr. Mallinson, who will be giving us a talk about his reminiscences in the field of magnetic recording over the last 40 years.

(Laughter)

DR. HARIHARAN: Dr. Mallinson has an M.A. Degree in Natural Philosophy in Physics from University College in Oxford, England. And he joined the Ampex Company in Harrisburg, Pennsylvania in 1954 to work on the theory and design of magnetic lodging elements.

In 1962, he joined Ampex Corporation in Redwood City, California, where he held many positions concerned with the understanding and development of magnetic recording systems. From 1976 to 1978, as Manager of Hybrid Magnetic Recording in the Data Systems Division, he was concerned with the initial design of the 750 MBS digital recorder.

From 1978 to 1984, he supervised the Magnetic Recording Technology Department Multidisciplinary Group, working in magnetic recording theory, high density head fabrication, coding and communication theory, and the exploration of advanced concepts in various areas of recording. In 1984, he was appointed as the Founding Director of the Center for Magnetic Recording Research at the University of California, San Diego. Since 1990, he has been the President of Mallinson Magnetics, Inc. He has published over 60 papers on a wide variety of topics in magnetic recording.

Dr. Mallinson was an IEEE Magnetic Society Distinguished Lecturer in 1983. In 1984, he was awarded the Alexander M. Poniatoff award, named after the founder of Ampex Corporation, an award for leadership in the theory and practice of magnetic recording. Dr. Mallinson?

(Applause)

BANQUET PRESENTATION

Professor John C. Mallinson

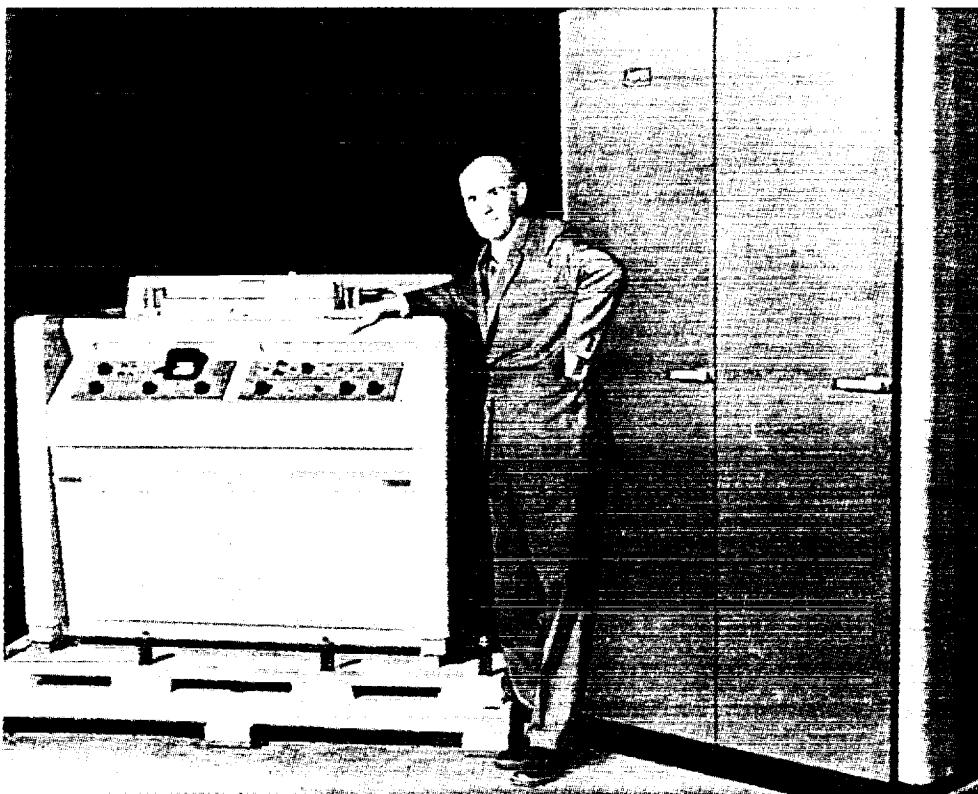
DR. MALLINSON: Thank you, Hari I've given a great deal of thought to what I should say to you this evening. Options ranged from telling you a number of risque jokes; but I realize that I have, in fact, been invited to give some reminiscences on--it's not a 40-year career; I'm not that old --

(Laughter)

DR. MALLINSON: But I have had a career since 1962 in magnetic recording. The most important of those years, I think, were the years at the Ampex Corporation; I was there for 24 years.

So, basically, what I'm going to talk about is some reflections on 24 years at Ampex.

First of all, I want to tell you that I think in the field of magnetic recording research and development, Ampex Corporation has very, very few equals. I won't list for you all the things that Ampex has developed; but most prominent amongst them is surely the invention of video tape recording.



ORIGINAL PAGE  
BLACK AND WHITE PHOTOGRAPH

DR. MALLINSON: There is a picture of Alex M. Poniatoff, the founder of Ampex, standing beside a video recorder. Alex M. Poniatoff hired me in 1962 to join the Ampex Corporation, and we hit it off together almost immediately because I was a pilot in the British Royal Air Force and still am a pilot. Alex M. Poniatoff was a pilot in the Russian Air Force in 1919.

And I will recount for you one story of Alex M. Poniatoff, and you should just think about this story while looking at Poniatoff when he is 76 years of age there.

In 1919, Alex M. Poniatoff was a Captain in the Russian Air Force. He was flying a wooden six-engined aircraft; I'm not sure what it was. Let's call it a Sigorsky; and they were flying it off some lake somewhere in Russia.

And one day, he went out to the lake and noticed that the red flag was flying; and he knew, as a good Captain in the Russian Air Force, that that meant that the commanding officer had decided that today the water was too rough for flying.

But Alex M. Poniatoff thought otherwise, as a young "Turk," as they say, in the Russian Air Force. He gathered up his copilot and his engineer, and they commanded someone to row them out to the flying boat; and they fired up the engines, turned it into the wind, and commenced the take-off run.

But I am not telling this story anywhere near as funny as he told the story; but long before the thing even got on the stack, before it even started showing any signs of becoming airborne, this wooden flying boat disintegrated. It broke up.

And shortly thereafter, he found himself in this icy cold water, calling out for his copilot and his engineer; and they all swam to each other and must have thought: What the hell are we going to do now?

And they noticed another small boat; this time it was a motor boat coming out, with their commanding officer in it. And he told me this hilarious story about what the commanding officer had to say about him; it was the end of his career flying the flying boats in the Russian Air Force.

So, that is Alex M. Poniatoff standing in front of the first Ampex video recorder in 1956. The video recorder is that unit; that unit weighed 1,100 pounds. And behind him is this cabinet which had 275 vacuum tubes in it; that was the electronics of the thing.

A reel of tape on that machine was 12 inches in diameter, 2 inch tape, and it recorded black and white television--NTSC television--for one hour. That was in 1956.

Before 1970, people had decided to use this video recorder for recording digital information; and such a reel of tape held 30 times  $10^9$  bits of information. I'll try and stick with bits all the time. Video people talk bits; computer people call big bits bytes.

They had 30 times  $10^9$ , 30 gigabytes, of data on it. And Ampex at that time, as was mentioned in the session today, started selling a system called the Terabit Memory, the  $10^{12}$  bit memory. A  $10^{12}$  bit memory had no less than 32 of those units.

And basically, what I'm going to talk about-- after some remarks--is how the evolution of recording since, say, 1960 to 1990, over the last 30 years, has resulted in much more compact ways of recording  $10^{12}$  bits of information.

At that time at Ampex, when these terabit systems were being made, with their 32 transports and their four Philco computers controlling them, we used to laugh at the idea of a terabit in a drawer; it was a joke. Well, a terabit in a drawer is now something which is available to everyone. I will show you a foil of it later on.

Rather than show foils all the time, which would just make it into a technical talk, I thought I would just plain talk to you about one of my major reflections at Ampex; and that is how it is that every two or three, perhaps every five, years it was perceived, though, that there was some important threat to the dominance of magnetic recording. And I'm afraid to tell you that that still goes on today, this notion that there is something around that is going to displace magnetic recording.

After 30 or 40 years of it, I very, very much doubt it, but I am getting ahead of myself.

The first thing that came along in the early 1960s, probably because so many of the Ampex engineers were involved in television and television revolves around cathode ray tubes-TV tubes--the notion came up that an electron beam recorder would be the best thing. It would be the answer to the maiden's prayer.

And Ampex, in the early 1960s, made electron beam recorders. The electron beam recorders recorded on photographic film that was specially made for us by Kodak; the track following servo was done by wobbling the electron beam, the dithering method of track following and observing scintillations of the scintillator coating on the back of the photographic film.

It recorded 100 megahertz analog bandwidth, which supported two 30 megabit per second channels. It was considered to be something wonderful in 1960 because the terabit memory was only about 5 megabits per second per channel, and this electron beam recorder had two 30 megabit per second channels.

What was wrong with the electron beam recorder like that? First of all, everything had to be done in a vacuum. Electron beams don't go too far in the air.

Secondly, it used photographic film, which had to be taken out of the vacuum and developed and put back in the vacuum to play back.

And thirdly, it was an enormously large machine. The electron column on it looked like a regular electron microscope. It was about 8 feet tall; and in fact, the customer for the electron beam recorder, the U.S. Air Force, found it necessary to increase the size of the cargo hatch of the C-130 at the time. It just wouldn't go in.

So, that was one of the things that had us all abuzz. Electron beam recording was going to solve all the problems. In retrospect, it's hard to know why people were so naive to think such a thing as that.

Shortly thereafter, a company was started in Boston; one of the principals, Dennis Speliotis, is here. It was called Macrobit. Macrobit had a similar notion that they were going to use electron beams to record on silicon wafers. Single crystal silicon had more or less just become available in the mid-1960s, and their notion was basically to lay out photolithographically little capacitors on a silicon wafer--tell me if I'm wrong, Dennis.

They were going to lay out little capacitors on a one-inch square, or maybe a one-and-a-half inch square, single silicon wafer; and it would have a capacity of about 10 megabits. And that one failed, I believe, because 10 megabits is a ridiculously small capacity.

At least at Ampex, it was realized that since Ampex's principal business was recording massive amounts of analog signals or data, 10 megabits was just too small a module to be of interest.

So, electron beam recording on silicon disappeared.

There was also a slight cautionary tale about that microbit memory in that its capacity did not exactly match the then-IBM disk drives. The IBM disk drives of the period were 3330s; and the Macrorbit memory had something like 7 percent more capacity than a single 3330, which in some spheres was its death knell because it's one thing to say that we will have an excess capacity in some memory device; and if we don't need to use all of that, we will just fill it in with garbage--you know, we'll just add noise, random 1's and 0's.

It is one thing to say that; it is altogether another thing to find some software engineers or some people who will do that. So, there's another lesson to be learned.

The two main reasons for electron beam recording were, of course, that you could do the recording without touching the medium. Touching the medium or not touching the medium was considered to be an extremely important fact; and that, in turn, led to the idea that there should be optical recorders of one kind or another.

I hate to think how much money the Ampex Corporation dumped into optical recording. Magneto-optical recording, which has been talked about this morning by the man from Alphatronix, is in its third--what's the right word?--life at the moment.

Originally, in the 1960s, the first life--the first period--of magneto-optical recording was being done on manganese bismuth; that was the material. You had to put in an awful lot of power on the manganese bismuth to heat it up to its Curie temperature of 400<sup>0</sup> or 350<sup>0</sup> centigrade.

And then, the material basically failed because, after repeated cycles to that high temperature, it changed phase; and the Curie temperature changed. It was unstable material; it would not stand indefinite recycling from room temperatures up to 400<sup>0</sup> centigrade.

Now, the second phase of magneto-optic recording led by Big Blue, was of course the Europium oxide run, which had terrible trouble that its compensation temperature was at liquid nitrogen temperatures. It failed. And the current one, which is iron cobalt--some rare earth--usually terbium--may or may not make it. Mr. Freese this morning said it would make it in a "niche" market; and I think that's about the right way to think about it.

What was the perception at Ampex of magneto-optic recording? Why, it was that Ampex should go ahead and make a magneto-optic tape recorder. And in 1966, there were shipped to the U.S. Navy no less than five magneto-optic tape recorders.

These tape recorders had 2 inch wide tape in them; the recording was transverse scan, just like the quad video machines. The transverse scan was done with a rotating polygonal mirror that was going round at 14,400 rpm, just the same speed as the quad drums. They were writing on a hard magnetic film with a permalloy overlayer and the readout scheme supported a data rate of 10 megabits per second.

We should have been smarter at that time to realize that this was in the mid-1960s; 10 megabits per second was already much slower than magnetic recorders were going in the mid-1960s.

But nevertheless, the hope persists in the human animal that there will be some other technology that will get you out of your perceived present problems.

The magneto-optic recorder was expensive; and I forgot to mention that, at that time in optical recording, there were no gallium arsenide solid-state diode lasers around. The lasers were all helium-neon gas lasers; and the gas lasers in this magneto-optic recorders shipped to the Navy were fully 24 inches long--helium-neon gas lasers. So, that failed.

I would say that basically the principal reason it failed was that it had material problems; and in particular, it had material problems where it was not possible to get more than 10 megabits per second data rate through it.

And that, I believe, is a persistent problem that is related with magneto-optic recording throughout all of its history including today. It is hard to get high data rates. So, it is a question to do with signal-to-noise ratio, which we don't have to go into.

50 megabits per second today is considered a high data rate for magneto-optic recording of any kind.

At the same time, or shortly after we realized that the magneto-optic recorder was not quite the right way to go, there was work undertaken in WORM, laser melting or ablating of some material like tantalum. It was very rapidly discovered that --

You know, the initial discovery phase of all of these endeavors is very exciting and goes very rapidly; and it is only later that we realize that you haven't got anything worthwhile.

In the case of the laser ablating, the pure WORM optical disk, it was very rapidly discovered that the laws of diffraction--Lord Rayleigh's laws of diffraction--do not apply in the writing phase. The writing phase is highly nonlinear. There is nothing that Lord Rayleigh ever said that precludes you from recording, say, one-tenth micron spots or 500 Angstrom spots.

And if you look back in the literature, in a journal called **SPIE**, the Society of Photographic and Instrumentation Engineers, you'll find papers from good old Ampex about recording 500 angstrom diameter holes which corresponded to an area density of nearly  $10^{10}$  bits to the square inch.

The hooker comes when you realize that what Rayleigh's diffraction limit--the area of the spot on the disk and all of that diffraction stuff--applies with a vengeance on playback. You can't ignore it on playback.

And the Airy disk is just like the gap loss function in magnetic recording. The only way to make the gap loss function in magnetic recording smaller or shorter is to make smaller gaps. The only way to make the area disk smaller in optical recording is either to think about numerical apertures that are ridiculous, more than unity, or go to shorter and shorter wavelengths.

At that time, using gas lasers, short wavelengths were available; and so, there were experiments done with blue lasers. And that possibility doesn't seem to be on at the moment with gallium arsenide. Gallium arsenide has an energy gap of about one electron volt, and that means it puts out photons that are around 8,000 Angstroms in wavelength--800 nanometers.

And there are not many semiconductor materials with larger energy gaps than that. A blue one at 4,000 Angstroms would need a 2 volt energy gap for instance.

So, we worked on WORMs. The basic trouble with a WORM is that it's a small module. A WORM disk of reasonable diameter, 6 inches or 8 inches in diameter, is only going to hold a gigabyte or so of data.

In the framework of Ampex thinking, that is nothing. That won't support digital video for very long at all. Digital video recording--the standards for it--were beginning to be set in the late 1970s; and Ampex, in fact, made the world's first digital video recorder.

It was a parallel access disk recorder, and it ran at 84 megabits per second. It was a composite video being sampled at four times the color subcarrier.

And so, it was realized that making a WORM disk, an optical disk, with perhaps  $10^{10}$  bits of information, wasn't going to support video for very long. So, that was canned.

The ROMs, the CD-ROMs, and the audio compact disk were never considered at Ampex because they are best regarded as a publication means. They are just a replication; it's a way of pressing disks to circulate information.

More or less at the same time, another pervasive idea in optical recording came up, which is still going to this day. There is a well-known laboratory in Texas that is still promoting the idea of doing holographic optical recording in three dimensions on crystals.

Ampex, of course, worked in that; the crystal, as always, was strontium niobate or lead niobate.

And I could go on for some time, if I wanted to make it a technical talk, about what's wrong with that idea. One of the things that is wrong with this holographic thing is that it depends on extreme mechanical precision in doing the optics.

It is not for nothing that when you go into an optic lab that you see everything is being done on granite blocks that are on air-bearing legs, and it's in an air-conditioned room. Holography works just fine, sending multiple beams through objectives. Lord knows in a telescope, or in this thing, whatever number you want to say, I'll agree with it. A million beams go through that lens.

But keeping them all in focus, keeping them all in the relative positions correctly, requires a dimensional stability and a vibration-free environment that makes it unlikely that it will work in any condition other than an optical bench in an air-conditioned room.

I'll tell you a little joke about optical recording because after-dinner speakers are supposed to tell jokes. It was told to me by a Dutchman, and the joke goes:

Do you know how you make a small fortune in optical recording? The answer is: You are either Edward Rothschild, who every year teaches courses with names like technology opportunity conference or something, or else you start with a large fortune --

(Laughter)

DR. MALLINSON: And that is exactly the way it has been in optical recording. Ask Schlumberger, ask Honeywell, ask Storage Technology.

So, what have I got left? I've got up there to about the mid-1970s; magneto-optic recording is coming around again in this third reincarnation in the oxide phase. Bubbles started to appear in the mid-1970s. Bubbles were, in my opinion, a loss leader right from the very beginning. The claim made by Andrew Bobeck of bubbles was that you would be able to get one million bubbles to the square inch and shift them at one million shifts per second.

Unfortunately, to do that, you had to have an entirely new material, a garnet substrate; but all my materials scientist friends told me it was going to be considerably more expensive to make than a silicon single crystal. That's overlooking the fact that single crystal silicon substrates were already available.

And then, if you looked at the structures that were involved in bubble recording, they were considerably more complicated; more area was required per bit cell. And bubbles have never gone anywhere.

There are still Japanese companies working on bubbles; I believe Hitachi still works on it. We at Ampex had a large bubble program because it was imagined, incorrectly, that bubbles were somehow going to do something for recording.

It's this constant notion that there is something there--some new technology--that hasn't been thought through very carefully; but it's new--new, new, new. New is the name of the game in research. That it would do something that semiconductors had not managed to do. I answered a question this morning about semiconductor memory; and believe it or not, good old, long-suffering Ampex had a semiconductor division down in Santa Monica in the L.A. Basin, and it was dedicated to trying to make a semiconductor memory.

And for those of you who weren't there this morning, the basic trouble with semiconductor memories is that there seems to be no chance that they are going to ever be economically worthwhile.

Magnetic recording, over the last 30 or 40 years, has continued to double in density every two or two and a half years.

I put forward the view this morning that the main reason that it was growing at that rate for such a long period of time was that that was the rate at which the electronics industry or the computer peripherals industry could accept technological change.

It just so happens that, apart from glitches to do with the dumping of semiconductors from you know where, the semiconductor industry has been doubling in density. They like to say every two years. In other words, it is following just about the same slope.

A gigabyte of semiconductor memory at the moment costs you \$500,000. A gigabit of magnetic memory in hard disk costs you \$1,000. So, there is a 500 to 1 differential in price; and the two technologies are advancing down the same sort of maturity curve.

If you take literally the slight differences in slope, then you come to the conclusion that they will reach equality in price in the year 2007. Now, that assumes that semiconductor technology can continue to advance in density at its current rate and so can magnetic.

In semiconductor technology, there seem to be some enormous barriers to do with good old diffraction of light coming up. Once you get down to quarter micron lines, it seems unlikely that even ultraviolet extreme blue light diffraction will do it. And you will have to go to X-ray lithography.

So, I don't believe that for a long time the semiconductor is going to be a threat to magnetic recording. It took Ampex a long time to work that out.

What in fact has happened, looking back with 20/20 hindsight--a retrospective look at things--is that every time there has been some advance in LSI, LSI has become larger, cheaper, denser, whatever the criterion is that you like to use--what has happened is that it has enabled some other form of recording to move to a higher level of performance.

At the Ampex Corporation, not only were rotary head video recorders invented in 1956, but helical scan recorders were invented two years later, in 1958. The helical scan recorder was invented because it had a long enough swipe down the length of the tape to get one whole television frame in--field--I beg your pardon.

Unfortunately, the time base errors of the head tape interface were much too high to be handled by the time base correctors that were in the early quad machines. The early quad machines--the very first ones--had mercury delay lines. The later ones had quartz crystals, where the delay time was changed by altering the voltage--the piezoelectric delay lines.

And that could not be used for the large timing errors in a long helical scan. Helical scan machines languished in the lab as an idle curiosity for almost five years, until, all of a sudden, I forget whether it was 8K or 16K DRAMs came along, that suddenly made it possible to do the time base correction digitally.

So, it was silicon technology that enabled helical machines to be useful. Likewise, in an optical disk at the moment--take the audio compact disk--the raw bit error rate coming off the disk, meeting the Philips or the manufacturing specification, is 1 byte--I shouldn't say byte; I promised I wouldn't say byte--let's say 1 bit error in 1,000, 1 in  $10^3$ , 1 in  $10^4$ . In order to get the error rate satisfactory for audio, which is 1 in  $10^{10}$ , which is one uncorrected bit in left/right stereo per hour. Large-scale integration error detection and correction Reed-Solomon and coding is used.

And just like in the RDAT, the rotary digital audio transport, so it is in the order of audio compact disk and the CD-ROM, that at any instant in time, no less than 64,000 bits of data are coming through the electronics, in transit, being corrected. This is doing the whole Reed-Solomon and coding business--you know, working out the syndrome, making the corrections, and all of that, doing those polynomial divisions.

And I would submit that there is another example: the audio compact disk, CD-ROM, the RDAT, rotary digital audio transport, are three machines which could not exist without large-scale integration of silicon.

While I was still working at the University of California at San Diego, people would repeatedly ask me-- the Chancellor and people like that; he was a psychologist of some kind; he can be forgiven for asking a question like this--but the question was repeated by other people, too. It is: When will silicon memory displace magnetic recording?

And I think the answer is clear from history. It will never do it. What will happen, as time goes on, is that the silicon memories that people use will get larger and larger and larger. If they are a megabit now, they will be 10 megabits in 1995 and 100 megabits in the year 2000.

Meanwhile, the recorder itself will have continued to increase in capacity. So, if floppy disks hold a megabit now, 1 megabit floppy disks won't exist ten years from now. They will all be silicon devices. Floppy disks will have moved up to 100 megabits and so on.

So, from the point of view of a recording person like me, I regard silicon technology as just being an enabling technology; it makes things possible.

So, that has finished, I think, all my discussion of the various threats that were thought. I've talked about electron beam recording of two kinds, optical recording of three kinds, including a magneto-optic tape recorder. I've talked about bubbles; I've talked about semiconductor memories, all of which must have cost Ampex an enormous amount of money and for which there is basically nothing to show.

And basically in the whole recording industry, there is almost nothing to show. It really is true--that sick joke--about starting off with a large fortune.

I want to just finish by telling you just a quick snapshot, just four foils about what has happened in recording.

(Showing of viewgraphs)

DR. MALLINSON: The most important thing that has happened since Mr. Poniatoff's day, 1956, is that the density of magnetic particles has increased enormously.

And on the top left here is a 1956 Ampex tape; and this is a 1966 Ampex tape, a 1976 tape, and a 1986 tape. It is pure g-Ferric oxide, more pure g-Ferric oxide, Cobalt-doped g-Ferric oxide, and metal particles. And there is a 1 micron marker down here, but you hardly need to see that. It's perfectly evident to you that something very, very dramatic has happened over that 30 year period in magnetic recording.

The density of particles in this iron particle tape is just 1,000 times higher than it was in the 1956 tape. You might ask: Well, why couldn't Ampex just jump to that immediately? The answer is that it is not making small particles that's difficult; it is formulating them, mixing them in the binder system, getting them uniformly dispersed--that's the difficult part. It seems to take a long time.

When you have a recording medium like that--and metal particle tape in my view, was the tape of the 1980s-- just as metal evaporated tape will be the tape of the 1990s--with a tape like that with 1,000 times more particles per unit volume, you get--if you are talking signal-to-noise power ratios--just 1,000 times the signal- to-noise ratio.

That means, track-width for track-width, wavelength for wavelength, 1,000 times the signal-to-noise power ratio, 30 dB signal-to-noise ratio.

(Change of viewgraph)

DR. MALLINSON: Ways of using that. For instance, the D-2 machine, to show you that it really exists; some of you heard discussion today about D-1 and D- 2. But there are two digital recording standards at the moment.

The D-2, which this is, takes the composite signal--that's the mixture of all the red-blue-green, if you will--and digitizes the whole thing at once; and a D-2 machine like this records that around 148 megabits per second.

The D-1 machine separates out the components--the red/blue/green, if you will--and digitizes each separately. And consequently, it requires a higher data rate; the data rate is 216 megabits per second.

So, there are two digital video recorders being made and in production in the world today: the D-1, which including overhead and eight audio channels and all of that, is 250 megabits per second; and the D-2 is 150 in this country and 160-something in Europe.

It's a rack-wired machine. The slot at the bottom takes cassettes.

(Change of viewgraph)

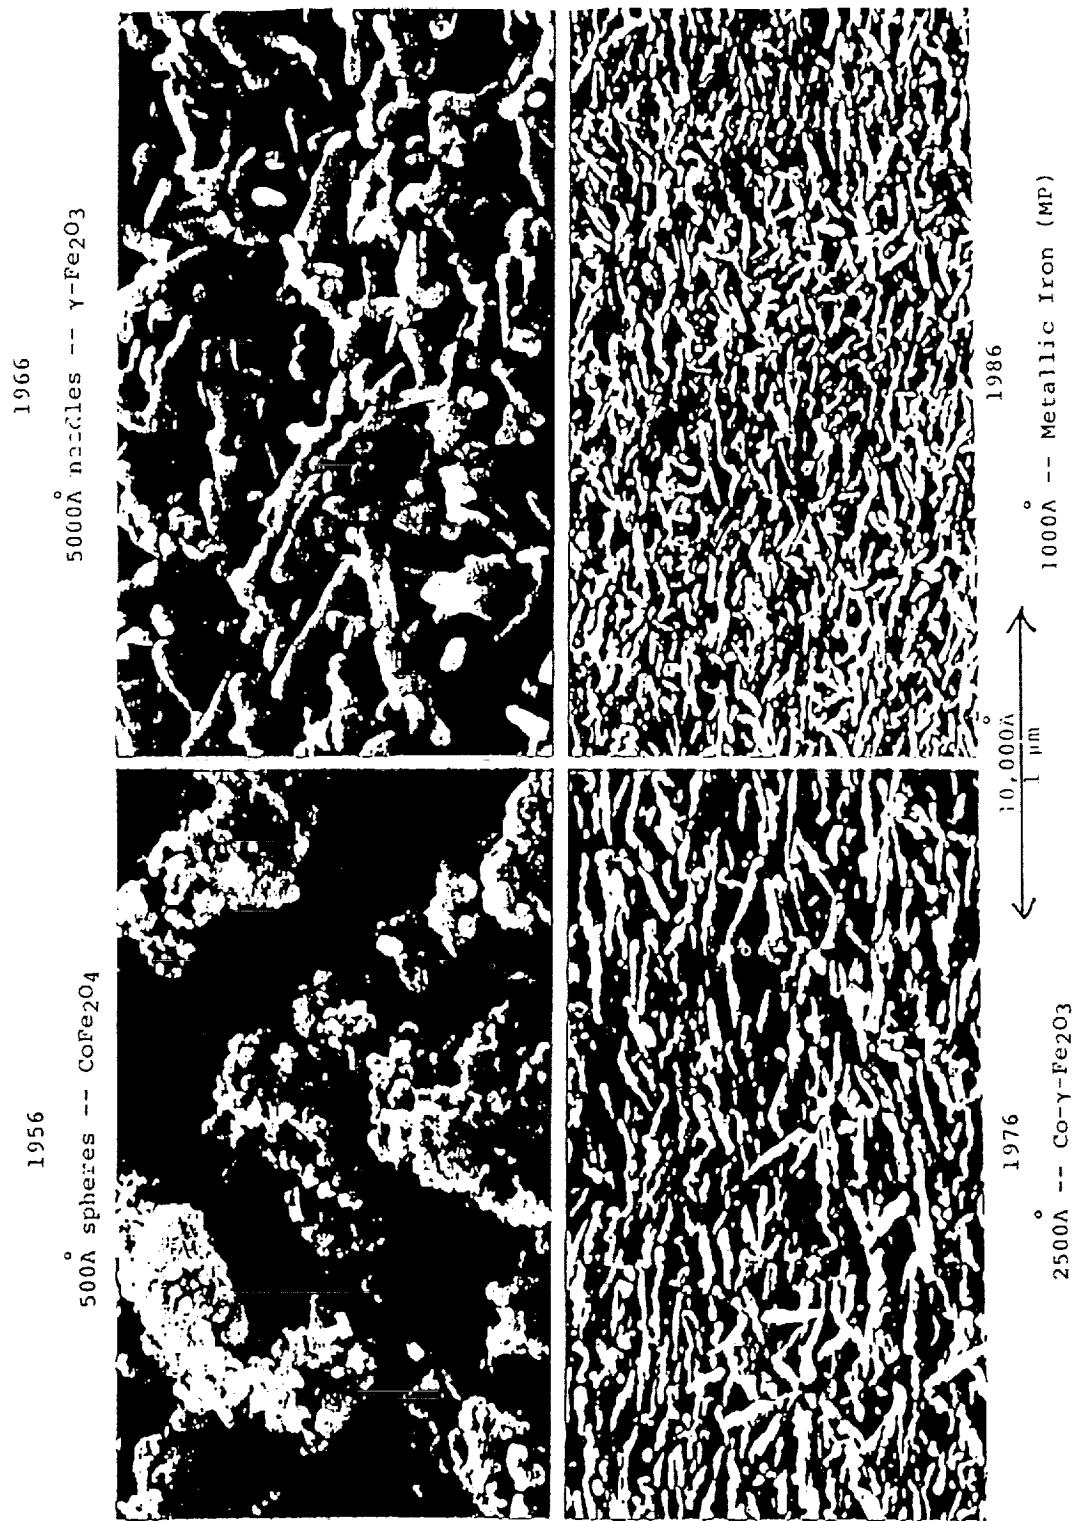
DR. MALLINSON: And the cassettes are called, not unreasonably, the large, the medium, and the small. I've got the dimensions on them there.

Let's look at the large one. It's 421 millimeters by whatever it is. It records for 3.5 hours-- 210 minutes--soaking up in Europe at 164 megabits per second. At the end of that 3.5 hours, it will have recorded 2 times  $10^{12}$  bits. So, there is 2 terabits there.

A cartridge like that is rather like two VHS cassettes put side by side. So, if Mr. Poniatoff was still around, I would tell him: Look, Alex, we really have got the TV end in a drawer now. Now, we have this drawer; and you can put in probably five or six of those. You can have close to  $10^{14}$  bits.

It is also interesting since we have been talking about alternative technologies to realize that one of those cartridges that costs about \$500.00--I'm sorry; they are called cassettes, not cartridges--has the same capacity as 30 IBM 3380 mainframe disk systems. And the mainframe disk systems cost \$100,000 apiece.

It is also equivalent in capacity to 40 two-sided 14 inch optical disks, which shows what I was saying earlier on about Ampex deciding that optical disks didn't have any future for a really high data rate or really large data storage is true.



That D-2 is, in fact, I believe--and if anyone disagrees me, I would like to hear--it is the largest digital store, single module digital store, available in the world today; it is 2 terabits of data.

(Change of viewgraph)

DR. MALLINSON: And just to show that I am not terribly prejudiced, here is another development in rotary head machines, television recorders. This is Sony's experimental digital high-definition TV recorder.

The simplest way to think about high-definition television is to think that it requires about five times the number of bits per frame as a regular television. Regular TV means about 3 megabits per frame; HDTV is about 30 megabits per frame.

Another way of thinking about HDTV is that it is rather like high resolution computer graphics that you might find on some work station, except the difference is that in TV you've got to put up an image, 30 of them per second. And the data rate with a recorder like this has to record is 1,100 megabits per second, 1.1 gigabits per second. And it is achieved; it is a rotary head machine. There is a rotary head hidden under there.

In this particular machine, which was made in 1984--it's not new technology by any means--it was achieved with six heads. So, each one of the heads contacting the tape was running at something like 180 or 190 megabits per second.

When I left Ampex in 1984, it was considered conservative practice in the design of machines to run at 80 megabits per second. This is higher--190. If you ask me what's the world's record at the moment in the published literature for a magnetic head, actually writing and playing back, it is a mind-blowing 300 megabits per second. So, I think Ampex was right in its decision; there was no other technology that was going to achieve such high data rates as magnetic recording.

A recorder like this records HDTV at 1,100 megabits per second for one hour; and the capacity, when you've finished, of these 12 inch reels is--it says at the bottom--4.5 terabits, which is equivalent to 100 two-sided 14 inch optical disks.

The race is on in all the recording companies these days to try and reduce the number of heads. It would be nice to get the number of heads down to four, for instance, not six.

So, that's enough for the foils.

The conclusion I want you to draw is that I told you the story about 25 years of R&D in the Ampex Corporation and a great deal of money and time and very skilled people's efforts were spent on activities which led nowhere. They did not lead to the goal of having extremely large databases to be accessed at extremely high data rates.

And I think that is the end of my talk. Thank you for your attention.

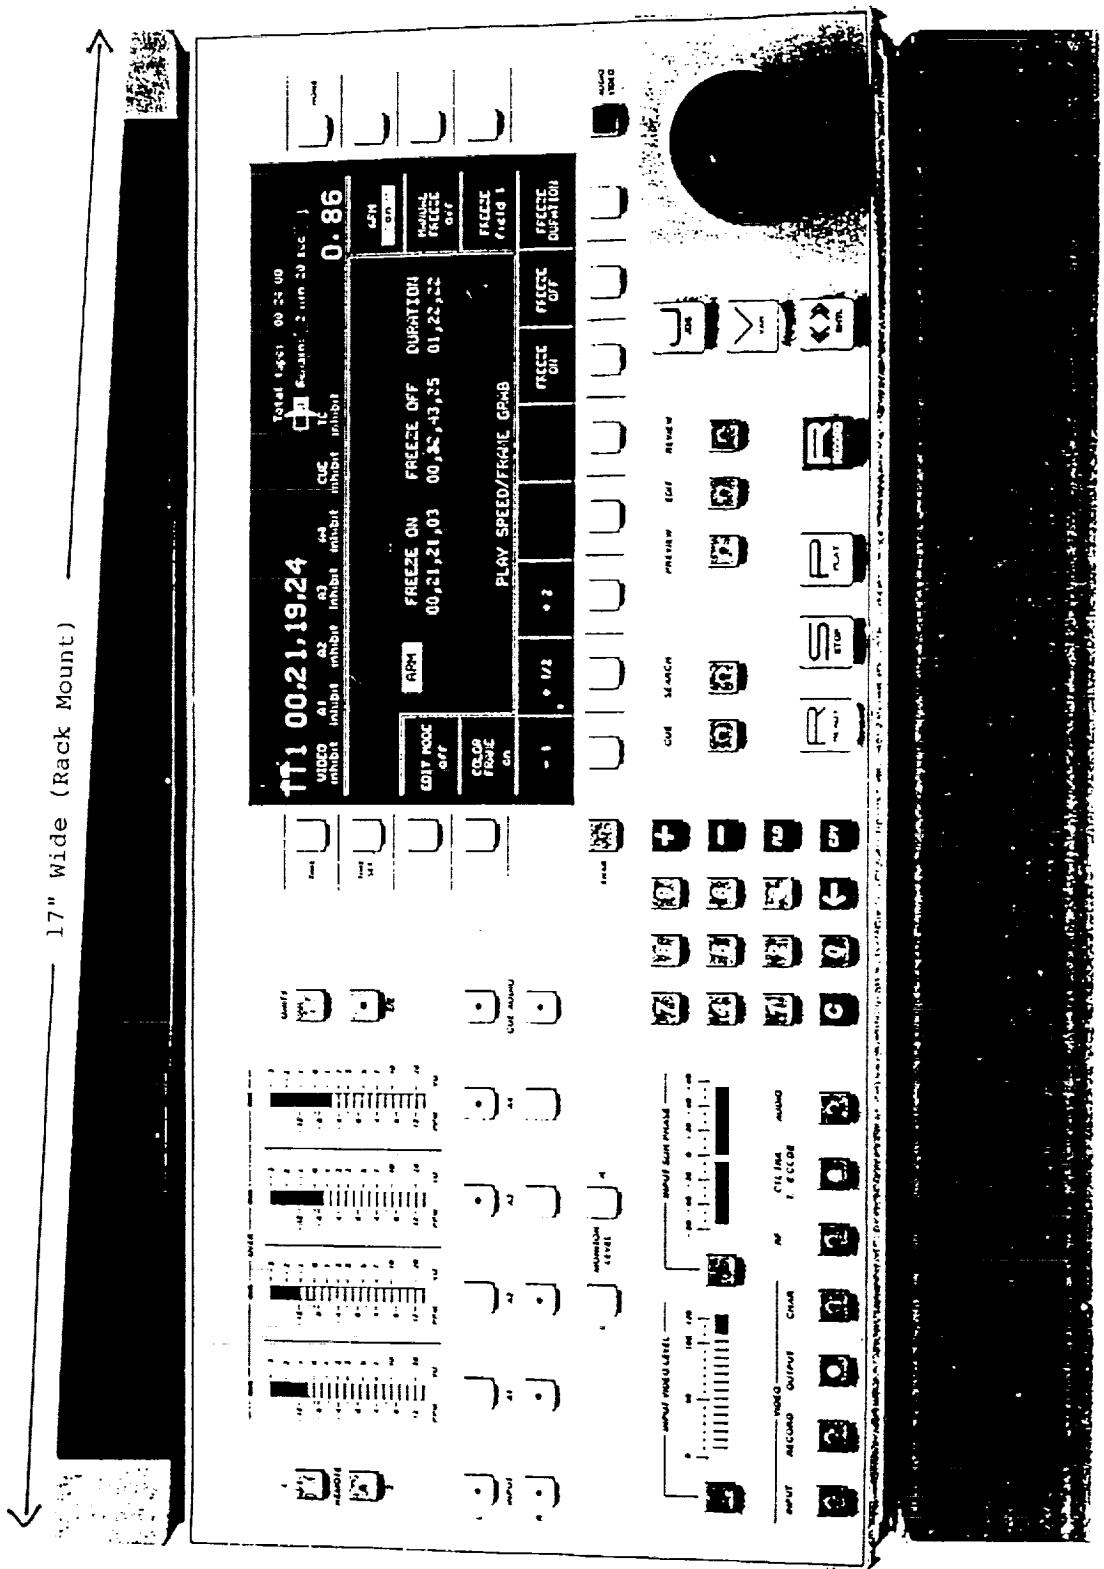
(Applause)

DR. MALLINSON: Questions?

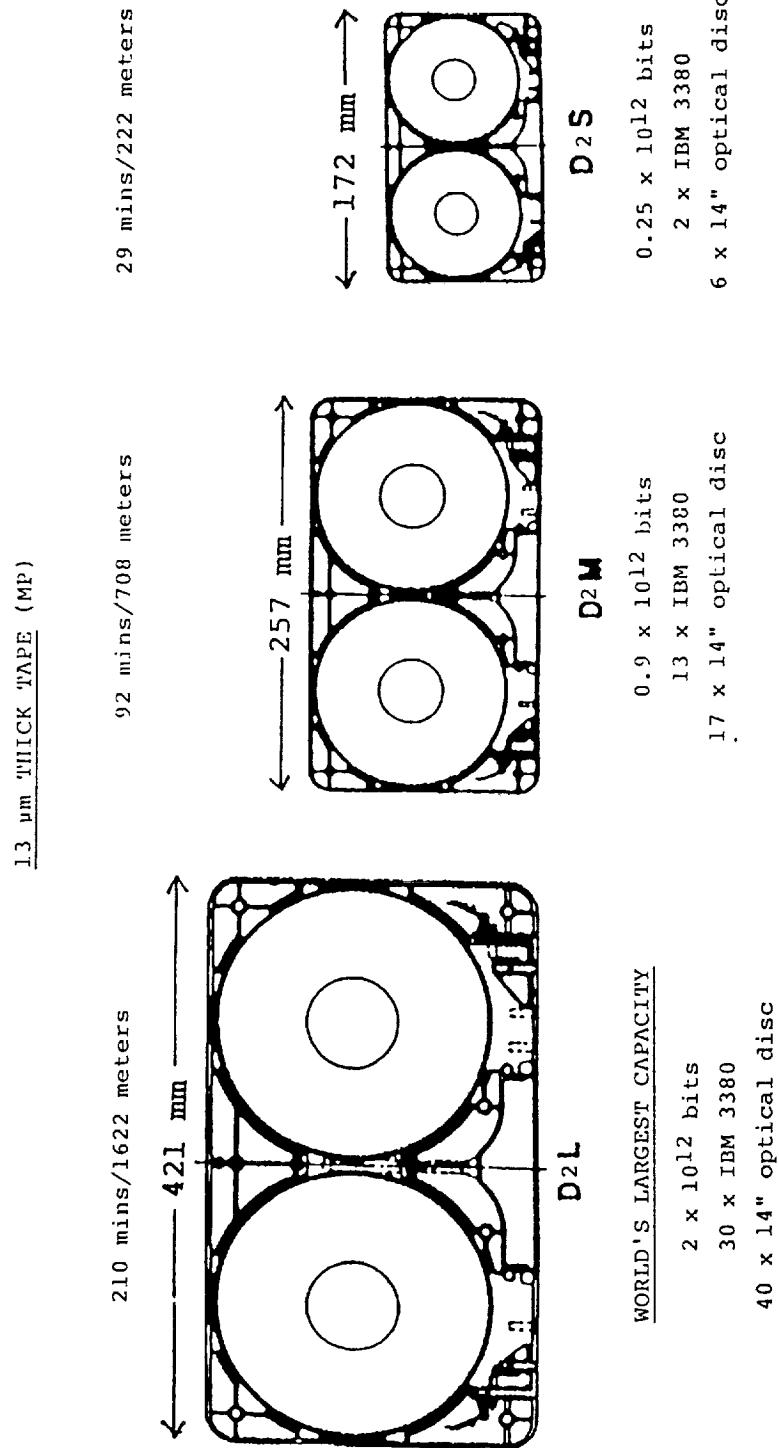
PARTICIPANT: Do you remember -- (Inaudible)

DR. MALLINSON: No.

PARTICIPANT: It was a holographic computer memory.

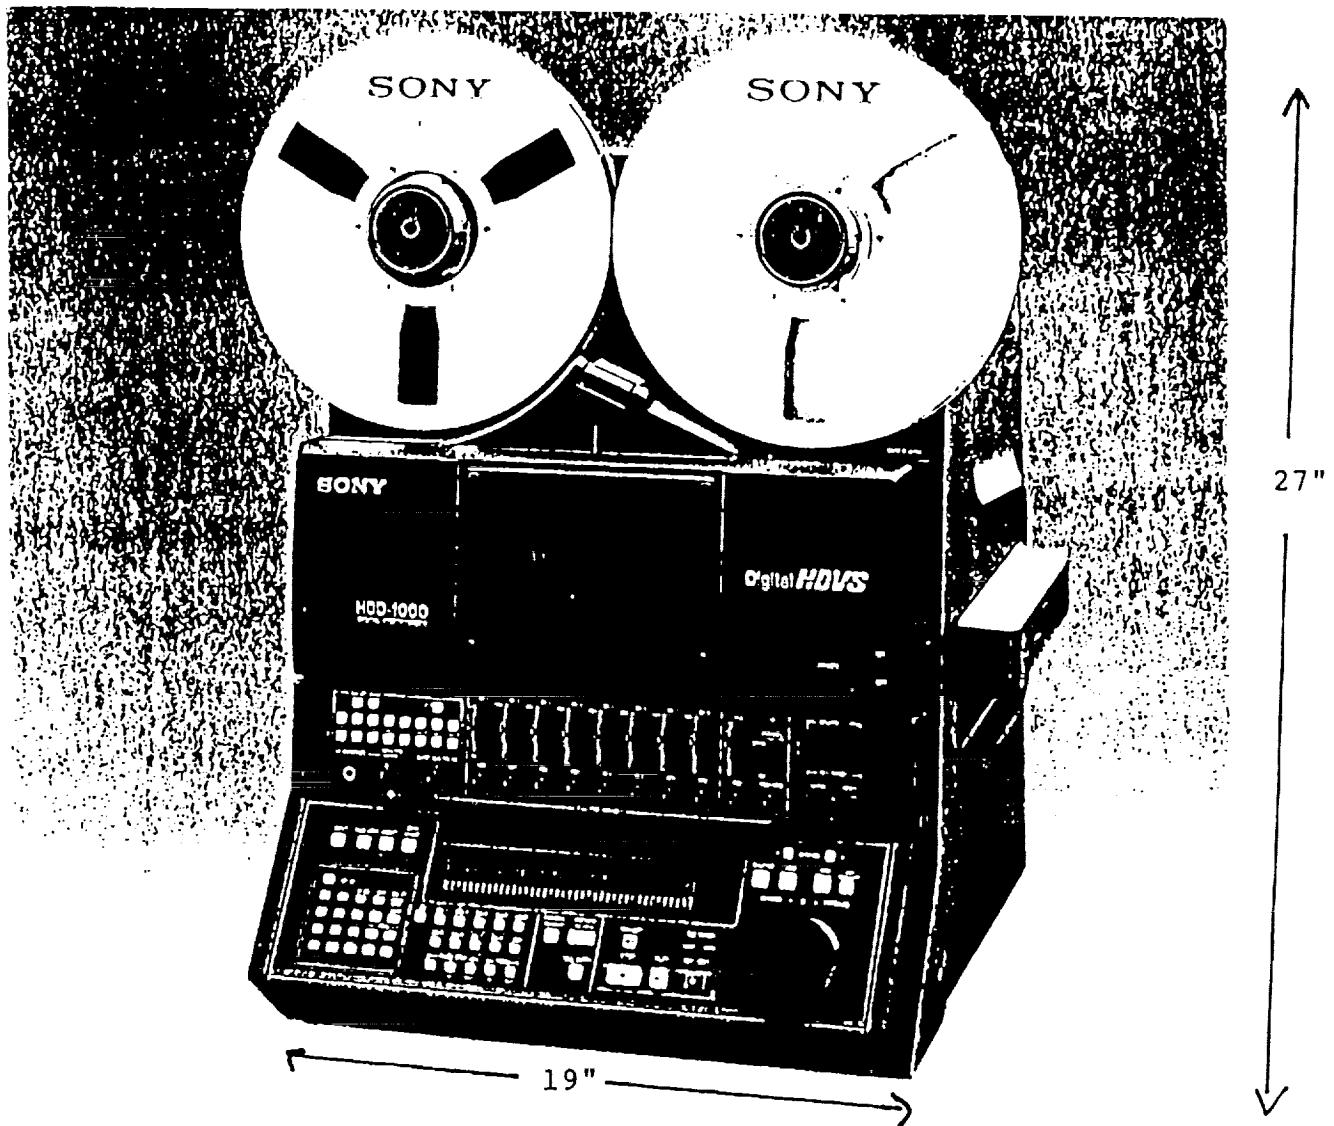


(D-2 Ampex VPR-300  
Digital Composite VTR)



#### 4.2.2 DIGITAL COMPONENT HDTV RECORDER

SONY HDD-1000 (Prototype only)



Sampling rates are:  $74:27:27 \times 10^6$  sample/sec

Video bit rate is 1100 + o/h Mbs

8 digital audio (48/16) channels also recorded (2 Mbs/channel)

Gross bit rate: about 1200 Mbs

Playing time: 63 minutes

Capacity:  $4.5 \times 10^{12}$  bits (100 x 14" optical discs)

DR. MALLINSON: No.

PARTICIPANT: (Inaudible)

DR. MALLINSON: It was a scan?

PARTICIPANT: Yes. It was in **Time** magazine. (Inaudible)

DR. MALLINSON: No, no. I don't recall that.

PARTICIPANT: (Inaudible)

DR. MALLINSON: You know, that has to limit the time one talks. I could talk for hours about this; I find it fascinating. And I would like to be able to say a few things about what you are supposed to learn from all this. Well, I suppose the things that you can learn from this is--and my remarks are just to do with tape recording--the goal of the Ampex Corporation is to record enormous databases, terabit databases, and access them at gigabit per second rates.

I think the message is that nothing will displace magnetic recording unless it is extremely simple; magnetic recording is done on a simple, chemically stable, featureless medium, it's cheap to make, it's cheap to implement.

I read all the excitement there is today about scanning tunneling microscopes; lo and behold, IBM can write IBM in some--do you remember what it was? Lithium or something? Xenon. It is evaporated away as soon as they let the temperature go up.

Then, the question is: Is there any reason to think that the very extreme high resolution of scanning tunneling microscopy--atomic resolution--will ever turn out to be a useful recording device? I think not.

Another one is atomic force microscopy, where you have a little magnetic tip that you measure the magnetic force over some substrate by vibrating it with a piezoelectric element and all of that. I think that one is even less likely to be a high density, fast access memory. In fact, I would add to the list of requirements that not only must the medium be cheap and featureless and stable, but there must be the requirement that you must be able to read every bit in at least 10 nanoseconds because 10 nanoseconds corresponds to 100 megabits per second.

And in fact, since all these magnetic recorders operate at that rate, and the 300 megabits per second I was talking about is only 3 nanoseconds to read each bit, a prerequirement really ought to be that you can read everything in just a few nanoseconds.

And then, there is the question of access time. Is there a way of accessing large distances? And is there a way of increasing the size of the recording medium to almost indefinitely large areas?

Questions like that, in atomic tunneling microscopy and atomic force microscopy, that I have found wanting.

PARTICIPANT: (Inaudible)

DR. MALLINSON: What I'm hearing is set size and tape height.

PARTICIPANT: Set size and tape height -- (Inaudible)

DR. MALLINSON: Cassette size?

PARTICIPANT: Yes. Real size -- (Inaudible)

PARTICIPANT: What's the difference in tension on the various sizes of cartridges that you displayed? And how does that affect -- (inaudible)

DR. MALLINSON: I hardly know how to answer. I mean, a standard tape is about a pound per inch of width, others are half-inch; and it is expected in video tape recorders where the tape contacts the head that the head shall be wearing down at the rate of about 10 microns per hour.

So, it's expected that the wear-out time of the head will be about 2,000 hours.

And a remark about cassettes may interest you. There is not a single cassette made--the D-1, the D-2, the 8 millimeter, the VHS, the late lamented Betamax--that even achieves 20 percent volume packing efficiency of the enclosure.

So, in order to have that convenience in magnetic recording in the cassette, which is simply the convenience that you can stick the thing in and pull it out at any time, you have given up a factor of 5 in volume packing density. 80 percent of the space inside the cassette is space, and they are still by far the highest volumetric packing factor devices in the world. With regard to another part of your question, I forgot to mention that it's funny that we used to be so worried about the tape contacting the head because any video tape now--even going by a VHS tape--to meet the specifications for VHS tape, it must stand one hour of still framing, which is something like 50,000 sequential passes of the head on the same track, without any measurable change.

And many of them will run for 10 hours.

PARTICIPANT: I have a question about tradeoff between areal density and the number of read -- (Inaudible)

DR. MALLINSON: There is nothing specifically about ID-1. I would be very surprised if any tape showed any deterioration in 100 reads. I'd even be surprised if it didn't go to 1,000 reads.

PARTICIPANT: (Inaudible)

DR. MALLINSON: I'm missing one word--the projection on the particle size?

PARTICIPANT: Yes.

PARTICIPANT: Could you repeat the question, please?

DR. MALLINSON: I think the question is: Do I have any idea of projections of the particle size? With iron particles that are in maybe RDAT and D-2 at the moment, they are around 1,200 Angstroms long and about 200 Angstroms wide. And it's interesting about that 200 Angstroms, incidentally, because I said this morning that various physicists have come up with the idea that there must be a quantum mechanical limit to recording density.

And my usual stock way is telling you the way to think about it some more is to tell them that in a single iron particle, 280 Angstroms in diameter, the flux flow is precisely one fluxon. Fluxons -- quantized magnetism has something to do with superconducting systems, not regular systems.

So, they are 1,200 by 200, and they will get smaller no doubt in the contact duplicating versions of DAT tape. The particle size is 800 angstroms long.

In the metallic thin films, the stuff I was talking about this morning, the metallic grains are 2 microinches by 2 microinches, 500 angstroms by 500.

There are media around at the moment which have a factor of 10 higher packing density than the iron particle tape. That's the name of the game with media; that's all you're doing. Raising the coercive force is a pain in the neck. It means more current has to go into the head; the medium must saturate; and who wants four times or five times the coercive force and 25 times the recording power?

The only reason you need it is to use smaller particles and get higher signal-to-noise ratios Any more questions?

(No response)

DR. MALLINSON: Thank you.

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188
<p>Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.</p>			
1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED	
	September 1992	Conference Publication, July 23-25, 1991	
4. TITLE AND SUBTITLE  NSSDC Conference on Mass Storage Systems and Technologies for Space and Earth Science Applications <i>Volume I</i>		5. FUNDING NUMBERS  JON-933-656-80-02-26	
6. AUTHOR(S)  Ben Kobler, P. C. Hariharan, and L. G. Blasso, Editors			
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)  NASA-Goddard Space Flight Center Greenbelt, Maryland 20771		8. PERFORMING ORGANIZATION REPORT NUMBER  92B00093 Code 933	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)  National Aeronautics and Space Administration Washington, D.C. 20546-0001		10. SPONSORING/MONITORING AGENCY REPORT NUMBER  NASA CP-3165, Vol. I	
11. SUPPLEMENTARY NOTES  Kobler: Goddard Space Flight Center, Greenbelt, MD; Hariharan and Blasso: STX Corporation*, Lanham, MD. *Hughes STX Corporation as of October 1, 1991			
12a. DISTRIBUTION/AVAILABILITY STATEMENT  Unclassified - Unlimited Subject Category 82		12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words)  This report contains copies of nearly all of the technical papers and viewgraphs presented at the NSSDC Conference on Mass Storage Systems and Technologies for Space and Earth Science Applications. This conference served as a broad forum for the discussion of a number of important issues in the field of mass storage systems. Topics include magnetic disk and tape technologies, optical disk and tape, software storage and file management systems, and experiences with the use of a large, distributed storage system. The technical presentations describe, among other things, integrated mass storage systems that are expected to be available commercially. Also included is a series of presentations from Federal Government organizations and research institutions covering their mass storage requirements for the 1990s.			
14. SUBJECT TERMS  Magnetic tape, magnetic disk, optical disk, mass storage, software storage		15. NUMBER OF PAGES 212	
		16. PRICE CODE A10	
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT UL